Rademacher Complexity Bounds for a Penalized Multiclass Semi-Supervised Algorithm

Yury Maximov Y.MAXIMOV@SKOLTECH.RU
Skolkovo Institute of Science and Technology
Skolkovo Innovation Center, Building 3-461
Moscow 143026 Russia

Massih-Reza Amini MASSIH-REZA.AMINI@IMAG.FR
University Grenobles Alps, CNRS, LIG
F38041 Grenoble Cedex 9 France

Zaid Harchaoui ZH13@NYU.EDU
Courant Institute for Mathematical Sciences,
715-719 Broadway, New York, NY 10003 USA

arXiv:1607.00567v2 [stat.ML] 29 Jul 2016

Abstract

We propose Rademacher complexity bounds for multiclass classifiers trained with a two-step semi-supervised model. In the first step, the algorithm partitions the partially labeled data and then identifies dense clusters containing κ predominant classes using the labeled training examples such that the proportion of their non-predominant classes is below a fixed threshold. In the second step, a classifier is trained by minimizing a margin empirical loss over the labeled training set and a penalization term measuring the disability of the learner to predict the κ predominant classes of the identified clusters. The resulting data-dependent generalization error bound involves the margin distribution of the classifier, the stability of the clustering technique used in the first step and Rademacher complexity terms corresponding to partially labeled training data. Our theoretical result exhibit convergence rates extending those proposed in the literature for the binary case, and experimental results on different multiclass classification problems show empirical evidence that supports the theory.

1. Introduction

Learning with partially labeled data, or Semi-supervised learning (SSL), has been an active field of study in the ML community these past twenty years. In this case, labeled examples are usually supposed to be very few leading to an inefficient supervised model, while unlabeled training examples contain valuable information on the prediction problem at hand which exploitation may lead to a performant prediction function. For this scenario, we assume available a set of labeled training examples $S_\ell = ((\mathbf{x}_i, y_i))_{i=1}^n \in (\mathcal{X} \times \mathcal{Y})^n$ drawn i.i.d. with respect to a fixed, but unknown, probability distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$ and a set of unlabeled training examples $S_u = (\mathbf{x}_i)_{i=n+1}^{n+u} \in \mathcal{X}^u$ supposed to be drawn from the marginal distribution, $\mathcal{D}_\mathcal{X}$, over the domain $\mathcal{X}$. If S_u is empty, then the problem is cast into the supervised learning framework. The other extreme case corresponds to the situation where S_ℓ is empty and for which the problem reduces to unsupervised learning.

The issue of learnability with partially labeled data was studied under three related yet different hypotheses of *smoothness assumption*, *cluster assumption*, and *low density separation* (Chapelle,

Schölkopf, Zien, et al., 2006; Zhu, 2005) and many advances have been made on both algorithmic and theoretical front under these settings.

Although classification problems, for which the design of SSL techniques is appealing, are multiclass in nature, the majority of theoretical results for semi-supervised learning has mainly considered the binary case (Balcan & Blum, 2010; El-Yaniv & Pechyony, 2009; Kääriäinen, 2005; Leskes, 2005; Urner, Shalev-Shwartz, & Ben-David, 2011). In this paper, we tackle the learning ability of multiclass classifiers trained on partially labeled data by first identifying dense clusters covering labeled and unlabeled examples and then minimizing an objective composed of the margin empirical loss of the classifier over the labeled training set, and also a penalization term measuring the disability of the learner to predict the predominant classes of dense clusters.

Our main result is a data-dependent generalization error bound for classifiers trained under this setting and which exhibits a complexity term depending on the effectiveness of the clustering technique to find homogenous regions of examples belonging to each class, the margin distribution of the classifiers and the Rademacher complexities of the class of functions in use defined for labeled and unlabeled data. The convergence rates deduced from the bound extends those proposed in the literature for the binary case, further experiments carried out on text and image classification problems, show that the proposed approach yields improved classification performance compared to extensions of state-of-the-art SSL algorithms to the multiclass classification case.

In the following section, we first define our framework, then the learning task we address. Section 3 presents the Rademacher generalization bound for a classifier trained with the proposed algorithm. Section 4 positions our theoretical findings with respect to the state-of-the-art, and finally, section 5 details experimental results that support this approach.

2. Penalized based semi-supervised multiclass classification

We are interested in the study of multiclass classification problems where the output space is $\mathcal{Y} = \{1, \ldots, K\}$, with $K > 2$. The semi-supervised multiclass classification algorithm that we consider is tailored under the cluster assumption and operates in two steps depicted in the following sections.

2.1 Partitioning of data and identifying κ-uniformly bounded clusters with level η

The first step consists in partitioning the labeled and the unlabeled training observations, Z, into $m > 0$ separate clusters with an algorithm $\mathcal{A}_Z$, $\Pi_Z = \{\mathcal{C}_1, \ldots, \mathcal{C}_m\}$.

Clusters of Π_Z that are well covered by classes in the labeled training set are then kept for learning the classifier (Section 2.2). Formally, for a fixed $\kappa \in \{1, \ldots, K\}$, let $\mathcal{Y}_\kappa(\mathcal{C})$ be the κ most predominant classes from $\mathcal{Y}$ present in cluster $\mathcal{C} \in \Pi_Z$. We then define κ-uniformly bounded clusters with level η, $\mathcal{C}_\kappa(\eta)$, the set of clusters within Π_Z that are covered by their κ most predominant classes such that the proportion of other classes within $\mathcal{C}$ not belonging to $\mathcal{Y}_\kappa(\mathcal{C})$ is less than η/m:

$$\mathcal{C}_\kappa(\eta) = \left\{ \mathcal{C} \in \Pi_Z : P_n\left((\mathbf{x}, y) \in S_\ell \cap \mathcal{C} \wedge y \notin \mathcal{Y}_\kappa(\mathcal{C})\right) \leq \frac{\eta}{m} \right\}. \tag{1}$$

Where P_n is the uniform probability distribution over S_ℓ; defined for any subset $B \subseteq S_\ell$, as $P_n(B) = \frac{1}{n}\text{card}(B)$.

2.2 Learning objective

In the second step, we address a learning problem that is to find, in a hypothesis set $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$, a scoring function $h \in \mathcal{H}$ with low risk:

$$R(h) = \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\left[\mathbb{1}_{g_h(\mathbf{x},y)\leqslant 0}\right], \tag{2}$$

where $\mathbb{1}_{\cdot}$ is the indicator function and $g_h(\mathbf{x}, y)$ is the margin of the function h at an example $(\mathbf{x}, y)$ (Koltchinskii & Panchenko, 2002):

$$g_h(\mathbf{x}, y) = h(\mathbf{x}, y) - \max_{y\neq y'} h(\mathbf{x}, y'). \tag{3}$$

This is achieved by minimizing a penalized empirical loss, defined for a given $\rho > 0$

$$\widehat{R}_\rho(h) = \widehat{R}_\rho(h, S_\ell) + \Omega_\rho(h, \mathcal{C}_\kappa(\eta)), \tag{4}$$

composed of an empirical margin loss of $h \in \mathcal{H}$ on a labeled training set S_ℓ,

$$\widehat{R}_\rho(h, S_\ell) = \frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell} \Phi_\rho(g_h(\mathbf{x}, y)), \tag{5}$$

and a penalization term that reflects the ability of the hypothesis $h \in \mathcal{H}$ to identify the κ most predominant classes within the disjoint clusters of $\mathcal{C}_\kappa(\eta)$;

$$\Omega_\rho(h, \mathcal{C}_\kappa(\eta)) = \frac{1}{u}\sum_{\mathcal{C}\in\mathcal{C}_\kappa(\eta)}\sum_{\mathbf{x}\in\mathcal{C}} \Phi_\rho(\rho_h(\mathbf{x}, \mathcal{Y}_\kappa(\mathcal{C}))), \tag{6}$$

where $\rho_h(\mathbf{x}, \mathcal{Y}_\kappa(\mathcal{C}))$ is the margin of an unlabeled example taken with respect to the set of κ predominant classes, $\mathcal{Y}_\kappa(\mathcal{C})$:

$$\rho_h(\mathbf{x}, \mathcal{Y}_\kappa(\mathcal{C})) = \max_{y\in\mathcal{Y}_\kappa(\mathcal{C})} h(\mathbf{x}, y) - \max_{y\in\mathcal{Y}\setminus\mathcal{Y}_\kappa(\mathcal{C})} h(\mathbf{x}, y),\ \mathbf{x} \in \mathcal{C} \subset \mathcal{C}_\kappa(\eta), \tag{7}$$

and, $\Phi_\rho : \mathbb{R} \to [0, 1]$ is the ρ-margin loss defined as (Koltchinskii & Panchenko, 2002) :

$$\forall z \in \mathbb{R}, \Phi_\rho(z) = \begin{cases} 0, & \text{if } \rho \leqslant z \\ 1 - z/\rho, & \text{if } 0 < z < \rho \\ 1, & \text{if } z \leqslant 0. \end{cases} \tag{8}$$

The pseudo-code of the proposed 2-step approach, referred to as Penalized Multiclass Semi-Supervised Learning ($\mathtt{PMS}_2\mathtt{L}$) in the following, is given in algorithm 1. The algorithm shares similarities with algorithms proposed in (Amini, Truong, & Goutte, 2008; Urner et al., 2011), where the k-NN technique was used to increase the size of the labeled training data by pseudo-labeling unlabeled examples that are in the nearest neighborhood of labeled examples, for binary classification and bipartite ranking. In (Rigollet, 2007), another two-step semi-supervised procedure is proposed where in the first stage a clustering of the feature space derived from the unlabeled data is produced and then each unlabeled observation, in a given cluster is assigned the same class label than the majority of labeled examples belonging to that class within the cluster. In the present

Algorithm 1: Pseudo-code of the PMS$_2$L algorithm

Input: Labeled data set $S_\ell = ((\mathbf{x}_i, y_i))_{i=1}^n \subseteq (\mathcal{X} \times \mathcal{Y})^n$;
Unlabeled data set $S_u = (\mathbf{x}_i)_{i=n+1}^{n+u} \subseteq \mathcal{X}^u$;
Hypothesis space $\mathcal{H}$;
m the number of clusters, $\mathcal{A}_{S_\ell \cup S_u} : \mathcal{X} \to \{1, \ldots, m\}$ a clustering algorithm found over $S_\ell \cup S_u$, $\kappa \in \mathbb{N}^*$, and $\eta > 0$;

Stage 1: Partition $S_\ell \cup S_u$ using $\mathcal{A}_{S_\ell \cup S_u}$; // The labels of examples in S_ℓ are not used
Using the labeled examples, S_ℓ, identify the κ-bounded clusters with level $\eta, \mathcal{C}_\kappa(\eta)$; // in accordance with Eq. (1)

Stage 2: Find a hypothesis $h^* \in \mathcal{H}$ that minimizes the penalized objective function (Eq. 4) :

$$h^* = \operatorname*{argmin}_{h \in \mathcal{H}} \widehat{R}_\rho(h)$$

Output: h^*

work we tackle a more general situation by considering multiclass classification problems and by relaxing the pseudo-labeling part which may be too aggressive in the multiclass case. Our analysis is based on the ability of a clustering technique to capture the structure of the data, and the ability of the classifier to identify predominant classes in κ-uniformly bounded clusters, leading to a multiclass definition of the cluster assumption which states that penalization over κ-uniformly bounded clusters with small level η helps learning.

3. Semi-supervised Data-dependent bounds

We now analyze how the use of unlabeled training data can improve generalization performance in some cases. Essentially, the trade-off is that clustering offers additional knowledge on the problem, therefore potentially helps learning, but can also be of lower quality, which may degrade it.

3.1 Notations and definitions

Before, let us first introduce notations that are used in the statement of the following results. We consider a hard clustering algorithm $\mathcal{A}_Z : \mathcal{X} \to \{1, \ldots, m\}$ defined as a function found over a finite sample Z such that each observation of $\mathcal{X}$ is assigned to one group label in the set $\{1, \ldots, m\}$.

Similarly to (Shamir & Tishby, 2007, 2008), we consider the distance between two clusterings $\mathcal{A}_Z$ and $\mathcal{A}_{Z'}$ defined over a set $\tilde{Z}$:

$$\Delta\left(\mathcal{A}_Z, \mathcal{A}_{Z'}, \tilde{Z}\right) = \min_\pi \sum_{\mathbf{x} \in \tilde{Z}} \mathbb{1}_{\mathcal{A}_Z(\mathbf{x}) \neq \pi(\mathcal{A}_{Z'}(\mathbf{x}))}. \tag{9}$$

Where $\pi : \{1, \ldots, m\} \to \{1, \ldots, m\}$ is a permutation. We suppose that the clustering algorithm $\mathcal{A}_\cdot$ obeys the bounded difference property (Shamir & Tishby, 2007; Von Luxburg, Belkin, & Bousquet, 2008). Namely, that there exists a constant L such that for any samples Z and Z' that differ in exactly one observation, we have $\Delta\left(\mathcal{A}_Z, \mathcal{A}_{Z'}, \tilde{Z}\right) \leqslant L$.

Note that the bounded difference property is tightly related to the notion of instability of a clustering method, and it has been studied for some clustering algorithms such as k-means or k-

hyperplane clustering. We refer to (Luxburg, 2010; Luxburg, Bousquet, & Belkin, 2004; Rakhlin & Caponnetto, 2006; Thiagarajan, Ramamurthy, & Spanias, 2011) and a number of references therein for the algorithmic details as well as various notions of clustering instability, and to (Shamir & Tishby, 2007) for the relation between stability, bounded differences property and model selection.

Further, the notion of function class capacity used here is the labeled and unlabeled Rademacher complexities of the function class $\mathcal{F}_{\mathcal{H}} = \{\mathbf{x} \mapsto h(\mathbf{x}, y) : y \in \mathcal{Y}, h \in \mathcal{H}\}$, defined respectively as:

$$\mathfrak{R}_n^*(\mathcal{F}_{\mathcal{H}}) = \sum_{\mathcal{C} \in \mathcal{C}_\kappa(\eta)} \mathbb{E}_{\sigma, S_\ell} \sup_{f \in \mathcal{F}_{\mathcal{H}}} \frac{2}{n} \left| \sum_{\mathbf{x}_i \in S_\ell \cap \mathcal{C}} \sigma_i f(\mathbf{x}_i) \right|,$$

$$\mathfrak{R}_u^*(\mathcal{F}_{\mathcal{H}}) = \sum_{\mathcal{C} \in \mathcal{C}_\kappa(\eta)} \mathbb{E}_{\sigma, S_u} \sup_{f \in \mathcal{F}_{\mathcal{H}}} \frac{2}{u} \left| \sum_{\mathbf{x}_i \in S_u \cap \mathcal{C}} \sigma_i f(\mathbf{x}_i) \right|,$$

$$\mathfrak{R}_n(\mathcal{F}_{\mathcal{H}}) = \mathbb{E}_{\sigma, S_\ell} \sup_{f \in \mathcal{F}_{\mathcal{H}}} \frac{2}{n} \left| \sum_{\mathbf{x}_i \in S_\ell \setminus \mathcal{C}_\kappa(\eta)} \sigma_i f(\mathbf{x}_i) \right|.$$

Finally, our results presented in the next section are based on the generalization bounds of the risks found over the κ-uniformly bounded clusters with level η. We refer to

$$R(h, \mathcal{C}_j) = \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mathbb{1}_{g_h(\mathbf{x},y)\leqslant 0} \wedge \mathbf{x} \in \mathcal{C}_j], \tag{10}$$

and corresponding

$$\widehat{R}_\rho(h, \mathcal{C}_j) = \frac{1}{n} \sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(g_h(\mathbf{x}, y)) + \frac{1}{u} \sum_{\mathbf{x}\in\mathcal{C}_j} \Phi_\rho(\rho_h(\mathbf{x}, \mathcal{Y}_\kappa(\mathcal{C}_j))), \tag{11}$$

as the risk and the empirical risk defined respectively over a single cluster $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$. Table 1 summarizes notations used throughout the paper.

Table 1: Notations

$\mathcal{X} \subseteq \mathbb{R}^d$	Input space,
$\mathcal{Y} = \{1, \ldots, K\}$	Output space,
m (resp. K)	Number of clusters (resp. classes),
S_ℓ (resp. S_u)	The set of labeled (resp. unlabeled) training examples,
$\mathcal{A}_Z : \mathcal{X} \to \{1, \ldots, m\}$	A clustering algorithm applied over the set Z,
$\mathcal{C}_\kappa(\eta)$	The set of κ-uniformly bounded clusters (Eq. 1),
$\mathcal{Y}_\kappa(\mathcal{C})$	κ most predominant classes found in cluster $\mathcal{C}$,
$g_h(\mathbf{x}, y)$	The margin of the function $h \in \mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$ at an example $(\mathbf{x}, y)$ (Eq. 3),
$\rho_h(\mathbf{x}, \mathcal{Y}_\kappa(\mathcal{C}))$	The margin of an unlabeled example taken with respect to $\mathcal{Y}_\kappa(\mathcal{C})$ (Eq. 7),
$\mu_h(\mathbf{x}) = \arg\max_{y\in\mathcal{Y}} h(\mathbf{x}, y)$	The class prediction of $h \in \mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$ for an example $\mathbf{x}$,
$\hat{R}_\rho(h)$	Penalized empirical loss (Eq. 4),
$\Omega_\rho(h, \mathcal{C}_\kappa(\eta))$	Penalization term in $\hat{R}_\rho(h)$ estimated over $\mathcal{C}_\kappa(\eta)$ (Eq. 6),
$\widehat{R}_\rho(h, \mathcal{C}_j)$	Empirical risk defined over a single cluster $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$,
P_n	Uniform probability distribution over S_ℓ.

3.2 Theoretical study

Theorem (2) is our main result and it sheds light on whether unlabeled data may help learning by providing bounds on the generalization error of a multiclass classifier trained with the penalized empirical loss stated above. It indicates situations where partitions, found by clustering, fit well to the natural structure of the data hence allowing to reach fast convergence rates.

The proof of the theorem is based on the next lemma that provides generalization bounds over the true risk estimated within a single confident cluster $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$ (Eq. 10).

Lemma 1 *Let $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$ be a hypothesis set where $\mathcal{Y} = \{1,\dots,K\}$, and let $S_\ell = ((\mathbf{x}_i, y_i))_{i=1}^n$ and $S_u = (\mathbf{x}_i)_{i=n+1}^{n+u}$ be two sets of labeled and unlabeled training data, drawn i.i.d. respectively according to a probability distribution over $\mathcal{X}\times\mathcal{Y}$ and a marginal distribution $\mathcal{D}_\mathcal{X}$. Fix $\rho > 0$, $\kappa \in \{1,\dots,K\}$ then for any $1 > \delta > 0$, the following multiclass classification generalization error bound holds with probability at least $1-\delta$ for all $h \in \mathcal{H}$ learned by the proposed algorithm over a single κ-uniformly bounded cluster $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$ derived from S_u by a clustering algorithm $\mathcal{A}_Z$ that obeys bounded differences property with constant $L \geqslant 1$:*

$$R(h,\mathcal{C}_j) \leq \widehat{R}_\rho(h,\mathcal{C}_j) + \frac{\eta}{m} + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}_\mathcal{H}) + \frac{2K}{\rho}\mathfrak{R}^*_{u,j}(\mathcal{F}_\mathcal{H}) + 3\sqrt{\frac{n^*_\eta(j)\kappa\log\frac{8K}{\delta}}{n^2}} + 3L\sqrt{\frac{u^*_\eta(j)\kappa\log\frac{8K}{\delta}}{u^2}},$$

where $n^*_\eta(j) = |S_\ell \cap \mathcal{C}_j|$, and $\mathfrak{R}^*_{n,j} = \mathbb{E}_{\sigma,S_\ell}\sup_{f\in\mathcal{F}_\mathcal{H}}\frac{2}{n}\left|\sum_{\mathbf{x}_i\in S_\ell\cap\mathcal{C}_j}\sigma_i f(\mathbf{x}_i)\right|$, and $u^*_\eta(j) = |S_u \cap \mathcal{C}_j|$, and $\mathfrak{R}^*_{u,j} = \mathbb{E}_{\sigma,S_u}\sup_{f\in\mathcal{F}_\mathcal{H}}\frac{2}{u}\left|\sum_{\mathbf{x}_i\in S_u\cap\mathcal{C}_j}\sigma_i f(\mathbf{x}_i)\right|$.

Proof. We start with the decomposition of the risk estimated in a single κ-uniformly bounded cluster $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$, by considering two situations where the prediction $\mu_h(\mathbf{x}) = \arg\max_{y\in\mathcal{Y}} h(\mathbf{x},y)$ falls within a set of confident clusters and without them respectively:

$$\begin{aligned} R(h,\mathcal{C}_j) = \mathbb{E}[\mu_h(\mathbf{x}) \neq y \wedge \mathbf{x}\in\mathcal{C}_j] \leqslant & \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x}) \neq y \wedge \mu_h(\mathbf{x}) = \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x}\in\mathcal{C}_j] + \\ & \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x}) \neq y \wedge \mu_h(\mathbf{x}) \neq \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x}\in\mathcal{C}_j] \quad (12) \end{aligned}$$

where $\mu_h(\mathbf{x},\mathcal{Y}_\kappa) = \arg\max_{y\in\mathcal{Y}_\kappa} h(\mathbf{x},y)$ and $\mathcal{Y}_\kappa = \mathcal{Y}_\kappa(\mathcal{C}_j)$ is a set of η–confident classes in $\mathcal{C}_j$.

The first term in the equality above is upper-bounded using the definition of the ρ-margin loss (Eq. 8)

$$\begin{aligned} \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x}) \neq y \wedge \mu_h(\mathbf{x}) = \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x}\in\mathcal{C}_j] &= \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x},\mathcal{Y}_\kappa) \neq y \wedge \mathbf{x}\in\mathcal{C}_j] \\ &\leqslant \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) \wedge \mathbf{x}\in\mathcal{C}_j], \end{aligned} \quad (13)$$

where $\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa) = h(\mathbf{x},y) - \max_{y'\in\mathcal{Y}_\kappa(\mathcal{C}_j)} h(\mathbf{x},y'), \mathbf{x}\in\mathcal{C}_j$.

Following (Lei, Dogan, Binder, & Kloft, 2015, remark 6), it comes out that for any fixed set $\mathcal{Y}'_\kappa \subset \mathcal{Y}$ of the predominant classes in cluster $\mathcal{C}_j$ we have with probability at least $1 - \delta/4K^\kappa$:

$$\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}'_\kappa)) \wedge \mathbf{x} \in \mathcal{C}_j] \leqslant$$
$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}'_\kappa)) + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}) + 3\sqrt{\frac{n^*_\eta(j)\log\frac{8K^\kappa}{\delta}}{n^2}} \leqslant$$
$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}'_\kappa)) + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}) + 3\sqrt{\frac{n^*_\eta(j)\kappa\log\frac{8K}{\delta}}{n^2}}, \tag{14}$$

Now for any possible set of κ predominant classes $\mathcal{Y}_\kappa$ in $\mathcal{C}_j$, and using the union bound and the inequality $\sum_{i=1}^k \binom{K}{i} \leqslant 2K^\kappa$, it comes from (14) that with probability at least $1 - \delta/2$ we have

$$\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) \wedge \mathbf{x} \in \mathcal{C}_j] \leqslant$$
$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}) + 3\sqrt{\frac{n^*_\eta(j)\kappa\log\frac{8K}{\delta}}{n^2}}, \tag{15}$$

By decomposing the sum in the first term of the above inequality, and considering the two cases where the class label y is within or without $\mathcal{Y}_\kappa$

$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) \leqslant \frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j \wedge y\in\mathcal{Y}_\kappa} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) + \frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j \wedge y\notin\mathcal{Y}_\kappa} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)),$$

Here we are in the case where $\mu_h(\mathbf{x}) = \mu_h(\mathbf{x},\mathcal{Y}_\kappa)$ (12) so, $\forall(\mathbf{x},y) \in S_\ell \cap \mathcal{C}_j \wedge y \in \mathcal{Y}_\kappa$, $\Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) = \Phi_\rho(g(\mathbf{x},y))$, and $\forall(\mathbf{x},y) \in S_\ell \cap \mathcal{C}_j \wedge y \notin \mathcal{Y}_\kappa, \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) \leqslant \mathbb{1}_{y\notin\mathcal{Y}_\kappa\wedge\mathbf{x}\in\mathcal{C}_j}$. Hence, for any sample S_ℓ and a set of predominant classes $\mathcal{Y}_\kappa$ we have

$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(\phi_h(\mathbf{x},y,\mathcal{Y}_\kappa)) \leqslant \frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(g_h(\mathbf{x},y)) + \frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell} \mathbb{1}_{y\notin\mathcal{Y}_\kappa\wedge\mathbf{x}\in\mathcal{C}_j},$$

From definition (1) we have $\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell} \mathbb{1}_{y\notin\mathcal{Y}_\kappa\wedge\mathbf{x}\in\mathcal{C}_j} \leqslant \eta/m$, and so

$$\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x}) \neq y \wedge \mu_h(\mathbf{x}) = \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x} \in \mathcal{C}_j] \leqslant$$
$$\frac{1}{n}\sum_{(\mathbf{x},y)\in S_\ell \cap \mathcal{C}_j} \Phi_\rho(g_h(\mathbf{x},y)) + \frac{\eta}{m} + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}) + 3\sqrt{\frac{n^*_\eta(j)\kappa\log\frac{8K}{\delta}}{n^2}}. \tag{16}$$

Futher, the second term in inequality (12) is upperbounded by

$$\begin{aligned}\mathbb{E}[\mu_h(\mathbf{x}) \neq y \wedge \mu_h(\mathbf{x}) \neq \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x} \in \mathcal{C}_j] &\leqslant \mathbb{E}_{\mathbf{x}\sim\mathcal{D}_\mathcal{X}}[\mu_h(\mathbf{x}) \neq \mu_h(\mathbf{x},\mathcal{Y}_\kappa) \wedge \mathbf{x} \in \mathcal{C}_j] \\ &\leqslant \mathbb{E}_{\mathbf{x}\sim\mathcal{D}_\mathcal{X}}[\Phi_\rho(\rho_h(\mathbf{x},\mathcal{Y}_\kappa)) \wedge \mathbf{x} \in \mathcal{C}_j],\end{aligned}$$

where $\rho_h(\mathbf{x},\mathcal{Y}_\kappa) = \max_{y\in\mathcal{Y}_\kappa(\mathcal{C}_j)} h(\mathbf{x},y) - \max_{y\in\mathcal{Y}\setminus\mathcal{Y}_\kappa(\mathcal{C}_j)} h(\mathbf{x},y)$, $\mathbf{x} \in \mathcal{C}_j$. As the clustering algorithm $\mathcal{A}_Z$ is supposed to obey the bounded differences property with constant L, then the function

$\mathcal{C}_j \mapsto \frac{1}{u}\sum_{\mathbf{x}\in\mathcal{C}_j}\Phi_\rho(\rho_h(\mathbf{x},\mathcal{Y}_\kappa))$ also obeys the independent bounded differences inequality with constant L (McDiarmid, 1989),(Ledoux & Talagrand, 1991, Theorem 4.12, p.112). Hence the standard Rademacher complexity bound (Bartlett & Mendelson, 2002, theorem 8) gives with probability at least $1-\delta/2$ uniformly for all $\mathcal{Y}_\kappa \in \mathcal{Y}$:

$$\begin{aligned}\mathbb{E}_{\mathbf{x}\sim\mathcal{D}_\mathcal{X}}[\Phi_\rho(\rho_h(\mathbf{x},\mathcal{Y}_\kappa)) \wedge \mathbf{x}\in\mathcal{C}_j] &\leqslant \frac{1}{u}\sum_{\mathbf{x}\in\mathcal{C}_j}\Phi_\rho(\rho_h(\mathbf{x},\mathcal{Y}_\kappa)) + \frac{2K}{\rho}\mathfrak{R}^*_{u,j}(\mathcal{F}) + 3L\sqrt{\frac{u^*_\eta(j)\log\frac{8K^\kappa}{\delta}}{u^2}} \\ &\leqslant \frac{1}{u}\sum_{\mathbf{x}\in\mathcal{C}_j}\Phi_\rho(\rho_h(\mathbf{x},\mathcal{Y}_\kappa)) + \frac{2K}{\rho}\mathfrak{R}^*_{u,j}(\mathcal{F}) + 3L\sqrt{\frac{u^*_\eta(j)\kappa\log\frac{8K}{\delta}}{u^2}},\end{aligned} \tag{17}$$

The statement of the Lemma follows from inequalities (12), (16) and (17). □

Our main result is a data-dependent generalization bound for any semi-supervised multiclass prediction function produced by the algorithm described above.

Theorem 2 *Let $\mathcal{H}\subseteq\mathbb{R}^{\mathcal{X}\times\mathcal{Y}}$ be a hypothesis set where $\mathcal{Y}=\{1,\ldots,K\}$, and let $S_\ell=((\mathbf{x}_i,y_i))_{i=1}^n$ and $S_u=(\mathbf{x}_i)_{i=n+1}^{n+u}$ be two sets of labeled and unlabeled training data, drawn i.i.d. respectively according to a probability distribution over $\mathcal{X}\times\mathcal{Y}$ and a marginal distribution $\mathcal{D}_\mathcal{X}$. Fix $\rho>0$, $\kappa\in\{1,\ldots,K\}$ then for any $1>\delta>0$, the following multiclass classification generalization error bound holds with probability at least $1-\delta$ for all $h\in\mathcal{H}$ learned by the proposed algorithm over the set of κ-uniformly bounded set of clusters, $\mathcal{C}_\kappa(\eta)$ derived from S_u:*

$$R(h) \leq \widehat{R}_\rho(h) + \eta + \frac{2K}{\rho}\mathfrak{R}^*_u(\mathcal{F}_\mathcal{H}) + \frac{2\kappa}{\rho}\mathfrak{R}^*_n(\mathcal{F}_\mathcal{H}) + \frac{2K}{\rho}\mathfrak{R}_n(\mathcal{F}_\mathcal{H}) + 3\sqrt{\frac{3\log\frac{12Km}{\delta}}{s_*}},$$

where $\frac{1}{s_*} \doteq L^2\frac{m\kappa u^*_\eta}{u^2} + \frac{m\kappa n^*_\eta}{n^2} + \frac{m(n-n^*_\eta)}{n^2}$, $n^*_\eta = |S_\ell\cap\mathcal{C}_\kappa(\eta)|$ and $u^*_\eta = |S_u\cap\mathcal{C}_\kappa(\eta)|$.

Proof. Let $\mathcal{C}=\{\mathcal{C}_1,\ldots,\mathcal{C}_m\}$ be a set of disjoint clusters derived from S_u. We decompose the risk of a classifier by considering the two exclusive cases whether the misclassification error occurs inside or outside the set of η-confident clusters.

$$\begin{aligned}R(h) = \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x})\neq y] = &\sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)}\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x})\neq y\wedge\mathbf{x}\in\mathcal{C}_j] + \\ &\sum_{\mathcal{C}_j\notin\mathcal{C}_\kappa(\eta)}\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x})\neq y\wedge\mathbf{x}\in\mathcal{C}_j]\end{aligned} \tag{18}$$

First, we bound the risk over the set of confident clusters. For any cluster $\mathcal{C}_j$ in $\mathcal{C}_\kappa(\eta)$ and any set of confident clusters $\mathcal{Y}_\kappa(\mathcal{C}_j)$ within it, from lemma 1 we have with probability at least $1-2\delta/3m$:

$$\begin{aligned}R(h,\mathcal{C}_j) = \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x})\neq y\wedge\mathbf{x}\in\mathcal{C}_j] \leqslant \widehat{R}_\rho(h,\mathcal{C}_j) + \frac{\eta}{m} + \frac{2\kappa}{\rho}\mathfrak{R}^*_{n,j}(\mathcal{F}_\mathcal{H}) + \frac{2K}{\rho}\mathfrak{R}^*_{u,j}(\mathcal{F}_\mathcal{H}) \\ + 3\sqrt{\frac{n^*_\eta(j)\kappa\log\frac{12Km}{\delta}}{n^2}} + 3L\sqrt{\frac{u^*_\eta(j)\kappa\log\frac{12Km}{\delta}}{u^2}},\end{aligned}$$

where $n_\eta^*(j) = |S_\ell \cap \mathcal{C}_j|$, and $\mathfrak{R}_{n,j}^*(\mathcal{F}) = \mathbb{E}_{\sigma,S_\ell} \sup_{f\in\mathcal{F}_\mathcal{H}} \frac{2}{n}\left|\sum_{\mathbf{x}_i\in S_\ell\cap\mathcal{C}_j} \sigma_i f(\mathbf{x}_i)\right|$, and $u_\eta^*(j) = |S_u \cap \mathcal{C}_j|$, and $\mathfrak{R}_{u,j}^*(\mathcal{F}) = \mathbb{E}_{\sigma,S_u} \sup_{f\in\mathcal{F}_\mathcal{H}} \frac{2}{u}\left|\sum_{\mathbf{x}_i\in S_u\cap\mathcal{C}_j} \sigma_i f(\mathbf{x}_i)\right|$.

Summing up over all clusters it comes

$$\sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} R(h,\mathcal{C}_j) \leq \sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} \widehat{R}_\rho(h,\mathcal{C}_j) + \eta + \frac{2\kappa}{\rho}\mathfrak{R}_n^*(\mathcal{F}_\mathcal{H}) + \frac{2K}{\rho}\mathfrak{R}_u^*(\mathcal{F}_\mathcal{H}) + 3\sum_{j=1}^m \sqrt{\frac{n_\eta^*(j)\kappa\log\frac{12Km}{\delta}}{n^2}} + 3L\sum_{j=1}^m \sqrt{\frac{u_\eta^*(j)\kappa\log\frac{12Km}{\delta}}{u^2}}.$$

By the Cauchy–Schwarz inequality $(\sum_{i=1}^m a_i b_i)^2 \leqslant (\sum_{i=1}^m a_i^2)(\sum_{i=1}^m b_i^2)$, then by fixing $b_i = 1, \forall i \in \{1,\ldots,m\}$, we can bound the two last terms of the right hand side inequality, and get

$$\sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} R(h,\mathcal{C}_j) \leq \sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} \widehat{R}_\rho(h,\mathcal{C}_j) + \eta + \frac{2\kappa}{\rho}\mathfrak{R}_n^*(\mathcal{F}_\mathcal{H}) + \frac{2K}{\rho}\mathfrak{R}_u^*(\mathcal{F}_\mathcal{H}) + 3\sqrt{\frac{mn_\eta^*\kappa\log\frac{12Km}{\delta}}{n^2}} + 3L\sqrt{\frac{mu_\eta^*\kappa\log\frac{12Km}{\delta}}{u^2}}, \tag{19}$$

with $n_\eta^* = \sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} = n_\eta^*(j)$ and $u_\eta^* = \sum_{\mathcal{C}_j\in\mathcal{C}_\kappa(\eta)} = u_\eta^*(j)$. Note, that a key property of a partition being part of the set of η–confident clusters $\mathcal{C}_\kappa(\eta)$ depends on the training set S_ℓ. Thus, in order to consider the data outside the set of confident clusters as i.i.d. we need to bound the risk over all possible separations of partitions as a set of confident cluster and its complementary. Then with probability at least $1-\delta/3$ the risk of classification outside the set of confident clusters is bounded as (Lei et al., 2015, remark 6)

$$\begin{aligned}
&\mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\mu_h(\mathbf{x}) \neq y \wedge \mathbf{x} \notin \mathcal{C}_\kappa(\eta)] \leqslant \\
&P_n(\mu_h(\mathbf{x}) \neq y \wedge \mathbf{x} \notin \mathcal{C}_\kappa(\eta)) + \frac{2K}{\rho}\mathfrak{R}_n(\mathcal{F}_\mathcal{H}) + 3\sqrt{\frac{(n-n_\eta^*)\log(6\frac{2m}{\delta})}{n^2}} \leqslant \\
&P_n(\mu_h(\mathbf{x}) \neq y \wedge \mathbf{x} \notin \mathcal{C}_\kappa(\eta)) + \frac{2K}{\rho}\mathfrak{R}_n(\mathcal{F}_\mathcal{H}) + 3\sqrt{\frac{(n-n_\eta^*)m\log(12/\delta)}{n^2}},
\end{aligned} \tag{20}$$

since the total number of labeled examples within confident clusters equals to n_η^*.

By the Cauchy-Schwartz inequality $(a+b+c)^2 \leqslant 3(a^2+b^2+c^2)$, we have

$$3\sqrt{\frac{mn_\eta^*\kappa\log\frac{12Km}{\delta}}{n^2}} + 3L\sqrt{\frac{mu_\eta^*\kappa\log\frac{12Km}{\delta}}{u^2}} + 3\sqrt{\frac{(n-n_\eta^*)m\log\frac{12}{\delta}}{n^2}} \leqslant 3\sqrt{3m\left(\frac{m\kappa n_\eta^*}{n^2} + L^2\frac{m\kappa u_\eta^*}{u^2} + \frac{m(n-n_\eta^*)}{n^2}\right)\log\frac{12Km}{\delta}} = 3\sqrt{\frac{3\log\frac{12Km}{\delta}}{s_*}}.$$

The result then follows from the inequalities (18), (19) and (20). □

We note that when clusters found by $\mathcal{A}$ cover well existing classes in S_ℓ, we have $\eta \approx 0$ from the definition (1). This condition hence guarantees that the generalization bound tends to zero when n and u tend to infinity. Hence, in such case unlabeled data generally bring valuable information for learning. Further, The data-dependent bound of theorem 2 can be explicitly written for certain class of functions, for example, for kernel-based hypotheses with $\mathfrak{K} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ a PDS kernel and $\Phi : \mathcal{X} \to \mathbb{H}$ its associated feature mapping function, defined as :

$$\mathcal{H}_B = \left\{ (\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y} \mapsto \langle \Phi(\mathbf{x}), \mathbf{w}_y \rangle \mid \mathbf{W} = (\mathbf{w}_1, \ldots, \mathbf{w}_K), \|\mathbf{W}\|_{\mathbb{H},2} \leq B \right\}.$$

Where $\|\mathbf{W}\|_{\mathbb{H},2}$ is the $L_{\mathbb{H},2}$ group norm of $\mathbf{W}$ defined as

$$\|\mathbf{W}\|_{\mathbb{H},2} = \sqrt{\sum_{k=1}^{K} \|\mathbf{w}_k\|_{\mathbb{H}}^2}.$$

In this case, we have the following corollary :

Corollary 3 *Let $\mathfrak{K} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ be a PDS kernel and let $\Phi : \mathcal{X} \to \mathbb{H}$ be the associated feature mapping function. Assume that there exists $R > 0$ such that $\mathfrak{K}(\mathbf{x}, \mathbf{x}) \leq R^2$ for all $\mathbf{x} \in \mathcal{X}$. Then for any $1 > \delta > 0$ and under the conditions and the definitions of theorem 2, the following multi-class classification error bound holds for all hypothesis $h \in \mathcal{H}_B$ learned by the proposed algorithm over the set of κ-uniformly bounded set of clusters, $\mathcal{C}_\kappa(\eta)$, with probability at least $1 - \delta$:*

$$R(h) \leq \widehat{R}_\rho(h) + \eta + \frac{4}{\rho} RB \sqrt{\frac{3k_*^2}{s_*}} + 3\sqrt{\frac{3 \log \frac{12Km}{\delta}}{s_*}}$$

where $\frac{1}{s_*} \doteq L^2 \frac{m\kappa u_\eta^*}{u^2} + \frac{m\kappa n_\eta^*}{n^2} + \frac{m(n-n_\eta^*)}{n^2}$ and $\frac{k_*^2}{s_*} \doteq K^2 \frac{mu_\eta^*}{u^2} + \kappa^2 \frac{mn_\eta^*}{n^2} + K^2 \frac{n-n_\eta^*}{n^2}$.

Proof. To proof the theorem we first bound the Rademacher complexities $\mathfrak{R}_u^*(\mathcal{F}_\mathcal{H})$, $\mathfrak{R}_n^*(\mathcal{F}_\mathcal{H})$ and $\mathfrak{R}_n(\mathcal{F}_\mathcal{H})$ and then apply theorem 2.

We follow the proposition (8.1) in (Mohri, Rostamizadeh, & Talwalkar, 2012) for bounding the Rademacher complexity of the class of linear classifiers in the feature space, and apply the Cauchy-Schwartz inequality $(\sum_{i=1}^m a_i b_i)^2 \leqslant (\sum_{i=1}^m a_i)^2 (\sum_{i=1}^m b_i)^2$ with $b_i = 1, \forall i$ and $a_i = \sqrt{u_\eta^*(i)}, \forall i$; to get

$$\mathfrak{R}_u^*(\mathcal{F}_\mathcal{H}) \leqslant \sum_{\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)} \frac{2}{u} RB \sqrt{u_\eta^*(j)} \leqslant 2RB \sqrt{\frac{mu_\eta^*}{u^2}},$$

where $u_\eta^*(j)$ in the number of unlabeled examples in η-confident cluster $\mathcal{C}_j$ and u_η^* is the total number of unlabeled examples within a set of confident clusters $\mathcal{C}_\kappa(\eta)$.

Similarly, if $n_\eta^*(j)$ is the number of unlabeled examples in $\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)$ we have

$$\mathfrak{R}_n^*(\mathcal{F}_\mathcal{H}) \leqslant \sum_{\mathcal{C}_j \in \mathcal{C}_\kappa(\eta)} \frac{2}{n} RB \sqrt{n_\eta^*(j)} \leqslant 2RB \sqrt{\frac{mn_\eta^*}{n^2}},$$

and by proposition 8.1 of (Mohri et al., 2012) we have $\mathfrak{R}_n(\mathcal{F}_\mathcal{H}) \leqslant 2RB\sqrt{\frac{n-n_\eta^*}{n^2}}$.

Applying the Cauchy-Schwartz inequality again we finally get

$$2K\mathfrak{R}_n(\mathcal{F}_\mathcal{H}) + 2\kappa\mathfrak{R}_n^*(\mathcal{F}_\mathcal{H}) + 2K\mathfrak{R}_u^*(\mathcal{F}_\mathcal{H}) \leqslant 4RB\sqrt{\frac{3k_*^2}{s_*}}$$

□

The non-empirical terms of this bound determine the convergence rate of the proposed penalized semi-supervised mutliclass algorithm, and hence following (Vapnik, 2000, theorem 2.1, p.38), gives insights on its consistency. These terms may be better explained using orders of magnitude (Knuth, 1976). If we now consider the common situation in semi-supervised learning where $u \gg n$, and that clusters found by $\mathcal{A}$ cover well existing classes in S_ℓ, leading to $\eta \approx 0, n_\eta^* \approx n, u_\eta^* \approx u$, and $\kappa = O(1)$, $L = O(1)$ and $m = O(K)$ then

$$\frac{k_*^2}{s_*} \doteq K^2\frac{mu_\eta^*}{u^2} + \kappa^2\frac{mn_\eta^*}{n^2} + K^2\frac{n-n_\eta^*}{n^2} = O\left(\frac{K^3}{u} + \frac{K}{n}\right)$$

and

$$\frac{1}{s_*} \doteq L^2\frac{m\kappa u_\eta^*}{u^2} + \frac{m\kappa n_\eta^*}{n^2} + \frac{m(n-n_\eta^*)}{n^2} = O\left(\frac{K}{u} + \frac{K}{n}\right)$$

the convergence rate of the bound of corollary 3 is of the order

$$\tilde{O}\left(\sqrt{\frac{K}{n}} + K\sqrt{\frac{K}{u}}\right), \tag{21}$$

where, for any real valued functions f and g the equality ; $f(z) = \tilde{O}(g(z))$ holds, if there exists a constant $\alpha > 0$ such that $f(z) = O(g(z)\log^\alpha g(z))$ (Knuth, 1976). In the following section we present an overview of the related-work and show that in the case where the clustering technique $\mathcal{A}$ captures the true structure of the data, measured by the set of κ-uniformly bounded clusters with rate η, resulting in approximations above, then for linear kernel-based hypotheses, the convergence rate (21) is the direct extension of dimension-free convergence rates proposed in semi-supervised learning for the binary case.

4. Related works and discussion

Semi-supervised learning (SSL) approaches exploit the geometry of data to learn a prediction function from partially labeled training sets (Seeger, 2000). The three main SSL techniques; namely graphical, generative and discriminant approaches, were mostly developed for the binary case and tailored under smoothness, low density separation and cluster assumptions (Chapelle et al., 2006; Zhu, 2005).

Graphical approaches construct an empirical graph where the nodes represent the training examples and the edges of the graph reflect the similarity between them. These approaches are mostly based on label spreading algorithms that propagate the class label of each labeled node to its neighbors (Zhou, Bousquet, Lal, Weston, & Schölkopf, 2003; Zhu, 2002).

Table 2: Summary of the convergence rates of dimension free bounds of excess risks for different SSL approaches.

Order of convergence rates	Case; Reference
$\tilde{O}\left(\frac{u\sqrt{\min(u,n)}}{n+u}\right)$	Binary; (El-Yaniv & Pechyony, 2009)
$\tilde{O}\left(\frac{1}{n}+\frac{1}{\sqrt{u}}\right)$	Binary;(Balcan & Blum, 2010)
$\tilde{O}\left(\frac{1}{\sqrt{n}}+\frac{1}{\sqrt{u}}\right)$	Binary; (Leskes, 2005)
$\tilde{O}\left(\frac{1}{\sqrt{n}}+\frac{1}{\sqrt{u}}\right)$	Binary; (Kääriäinen, 2005)
$\tilde{O}\left(\frac{\sqrt{K}}{\sqrt{n}}+\frac{K^{3/2}}{\sqrt{u}}\right)$	Multi-class; Corollary 3

Generative approaches naturally exploit the geometry of data by modelling their marginal distributions. These methods are developed under the cluster assumption and use the Bayes rule to make decision. In the seminal work of (Castelli & Cover, 1995) it is shown that, without extra assumptions relating marginal distribution and true distribution of labels, a sample of unlabeled data is of (almost) no help for learning purpose. Recent work from (Ben-David, Lu, & Pál, 2008) investigated further the limitations of semi-supervised learning and concluded that theoretical results for semi-supervised learning should be accompanied by an extra assumption on the true label distribution.

Discriminant approaches directly find the decision boundary without making any assumptions on the marginal distribution of examples. The two most popular discriminant models are without doubts co-training (Blum & Mitchell, 1998) and Transductive SVMs (Vapnik, 2000). The co-training algorithm supposes that each observation is produced by two sources of information and that each view-specific representation is rich enough to learn the parameters of the associated classifier in the case where there are enough labeled examples available. The two classifiers are first trained separately on the labeled data. A subset of unlabeled examples is then randomly drawn and pseudo-labeled by each of the classifiers. The estimated output by the first classifier becomes the desired output for the second classifier and reciprocally. Under this setting, (Leskes, 2005) proposed a Rademacher complexity bound, where unlabeled data are used to decrease the disagreement between hypotheses from a class of functions $\mathcal{H}$ and proved that in some cases, the bound of the excess risk $|R(h) - \hat{R}(h, S_\ell)|$ for any $h \in \mathcal{H}$ is of the order $\tilde{O}\left(n^{-1/2} + u^{-1/2}\right)$. Another study in this line of research is (Tolstikhin, Zhivotovskiy, & Blanchard, 2015). However, transductive learning tends to produce a prediction function for only a fixed number of unlabeled examples. Transductive algorithms generally use the distribution of unsigned margins of unlabeled examples in order to guide the search of a prediction function and find the hyperplane in a feature space that separates the best labeled examples and that does not pass through high density regions. The notion of transductive Rademacher complexity was introduced in (El-Yaniv & Pechyony, 2009). In the best case, the excess risk bound proposed in this paper is of the order $\tilde{O}\left(u\sqrt{\min(u,n)}/(n+u)\right)$.

Our two step multiclass SSL approach is in between generative and discriminant approaches, and hence bears similarity with the study of (Urner et al., 2011). The main difference is however that the proposed approach does not rely on any pseudo-labeling mechanism and that our analyzes are based on the Rademacher complexity leading to dimension free data-dependent bounds. On another level and under the PAC-Bayes setting, (Kääriäinen, 2005) showed that in the realizable case where the hypothesis set contains the Bayes classifier, the obtained excess risk bound takes the form $\inf_{f\in F_0}\sup_{g\in F_0}\hat{d}(f,g)+\tilde{O}\left(u^{-1/2}\right)$; where $\hat{d}(f,g)$ is a normalized empirical disagreements between two hypothesis that correctly classify the labeled set and can be of order at least $\tilde{O}\left(n^{-1/2}\right)$.

The convergence rates of the mentioned bounds are sum up in Table 2. From these results, it becomes apparent that the convergence rate deduced from corollary 3, (Equation 21) extends those found in (Kääriäinen, 2005; Leskes, 2005) and obtained under different other settings.

5. Experimental Results

We perform experiments on six publicly available datasets. The three first ones are `Fungus`, `Birds` and `Athletics` that consist of three aggregations of lead nodes that go down from parent nodes in the ImageNet hierarchy[1]. Each image is characterized by a Fisher vector representation as described in (Harchaoui, Douze, Paulin, Dudík, & Malick, 2012). The three others collections are respectively the `MNIST` database of handwritten digits, the pre-processed 20 Newsgroups (`20-NG`) collection[2] and the `USPS` dataset[3]. Table 2 resumes the characteristics of these datasets. The proportions of training and test sets were kept fixed to those given in the released data files. Within the training set ($S_\ell \cup S_u$) we randomly sampled labeled examples S_ℓ, with different sizes, and used the remaining as unlabeled data.

Table 2: Characteristics of datasets used in our experiments.

dataset	$\lvert S_\ell \cup S_u\rvert$	size of the test	dimension, d	# of classes, K
`Birds`	5785	5596	4096	196
`Athletics`	28752	28727	4096	51
`Fungus`	50270	50271	4096	134
`20-NG`	15936	3393	62061	20
`MNIST`	60000	10000	780	10
`USPS`	7291	2007	256	10

To validate the proposed penalized based multiclass semi-supervised learning approach (`PMS`$_2$`L`), we compared its results with respect to a multiclass extension of a popular SSL algorithm proposed within each of the Generative, Graphical and Discriminant approaches. More precisely we considered the extension of the label propagation algorithm to the multiclass case (`McLP`) proposed by (Wang, Tu, & Tsotsos, 2013). A generative SSL model based on the mixture of gaussians (`S`$_2$`GM`), the extension of TSVM[4] (Vapnik, 2000) to the multiclass case (`McTSVM`), and a purely supervised technique which does not make use of any unlabeled examples in the training stage (`SUP`).

1. `http://www.image-net.org/challenges/LSVRC/2010/`
2. `http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multiclass.html`
3. `http://www-i6.informatik.rwth-aachen.de/~keysers/usps.html`
4. `http://svmlight.joachims.org/`

As the clustering algorithm $\mathcal{A}$, we employed the Nearest Neighbor Clustering technique proposed in (Bubeck & Luxburg, 2009), and fixed $m = 4K$, $\kappa = 2$ and $\eta = 10^{-3}$. Meaning that each cluster in $\mathcal{C}_\kappa(\eta)$ is mainly composed of the two most predominant classes within it. For the second stage of PMS$_2$L, as well as for SUP and McTSVM, we adapted the aggregated one-versus-all approach using a linear kernel SVM that respects the conditions of corollary 3. The penalized objective function can be easily implemented using convex optimization tools for convex surrogates of the 0/1 loss. The parameter C of the SVM classifier is determined by five fold cross-validation in logarithmic range between 10^{-4} and 10^4 over the available labeled training data. Results are evaluated over the test set using the accuracy, and the reported performance is averaged over 25 random (labeled/unlabeled/test) sets of the initial collections.

Table 3 summarizes results obtained by SUP, PMS$_2$L, McLP, S$_2$GM and McTSVM when a very small proportion of labeled training data is used in the learning of the models. We use boldface to indicate the highest performance rates, and the symbol $^\downarrow$ indicates that performance is significantly worse than the best result, according to a Wilcoxon rank sum test used at a p-value threshold of 0.05 (Lehmann, 1975). From these results it becomes clear that

- The algorithm PMS$_2$L performs significantly better than all of the four other algorithms, and it improves over SUP by an average of 1.5 to 6.5% on different datasets.

- McLP and McTSVM also perform better than SUP, though not in the same range than previously, while the mixture of Gaussians S$_2$GM does worse than SUP especially in the cases where the dimension of the problem is high.

- Finally, the difference in performance between PMS$_2$L and McTSVM is smaller than the one between the former and McLP.

Our analysis of these results is that the Nearest Neighbor Clustering technique (Bubeck & Luxburg, 2009) is effectively able to map correctly the considered data, into homogenous clusters containing mostly unlabeled examples of the same class than the $\kappa = 2$ most predominant classes contained in them. In this case, the penalized term of the objective function used to learn the classifier (Equation 4) forcefully helps to pick a better hypothesis in the set of linear classifiers,

Table 3: Means and standard deviations of the classification accuracy on test data over the 25 trials for each data set. n_y refers to the average number of labeled examples per class in each data set. $^\downarrow$ indicates statistically significantly worse performance than the best result, shown in bold, according to a Wilcoxon rank sum test ($p < 0.05$) (Lehmann, 1975).

Dataset	n_y	$n/(n+u)$	SUP	PMS$_2$L	McLP	S$_2$GM	McTSVM
Birds	5	0.18	$.294^\downarrow_{\pm.03}$	$\mathbf{.344}_{\pm.03}$	$.303^\downarrow_{\pm.06}$	$.286^\downarrow_{\pm.08}$	$.312^\downarrow_{\pm.04}$
Athletics	43	0.08	$.258^\downarrow_{\pm.03}$	$\mathbf{.273}_{\pm.02}$	$.259^\downarrow_{\pm.05}$	$.246^\downarrow_{\pm.07}$	$.263_{\pm.04}$
Fungus	15	0.04	$.121^\downarrow_{\pm.03}$	$\mathbf{.160}_{\pm.03}$	$.125^\downarrow_{\pm.06}$	$.107^\downarrow_{\pm.05}$	$.134^\downarrow_{\pm.04}$
20-NG	16	0.02	$.468^\downarrow_{\pm.05}$	$\mathbf{.531}_{\pm.03}$	$.476^\downarrow_{\pm.06}$	$.452^\downarrow_{\pm.04}$	$.484^\downarrow_{\pm.04}$
MNIST	120	0.02	$.767^\downarrow_{\pm.03}$	$\mathbf{.799}_{\pm.02}$	$.771^\downarrow_{\pm.05}$	$.758^\downarrow_{\pm.06}$	$.781^\downarrow_{\pm.01}$
USPS	14	0.02	$.790^\downarrow_{\pm.03}$	$\mathbf{.821}_{\pm.02}$	$.796^\downarrow_{\pm.04}$	$.788^\downarrow_{\pm.06}$	$.801^\downarrow_{\pm.02}$

Figure 2: Accuracy in percentage with respect to the proportion of labeled examples in the initial training set for ImageNet `Birds` (a), `Athletics` (b), `Fungus` (c); `20-NG` (d), `MNIST` (e), and `USPS` (f). Each reported performance on the test is averaged over 25 random (labeled/unlabeled/test) sets of the initial collections.

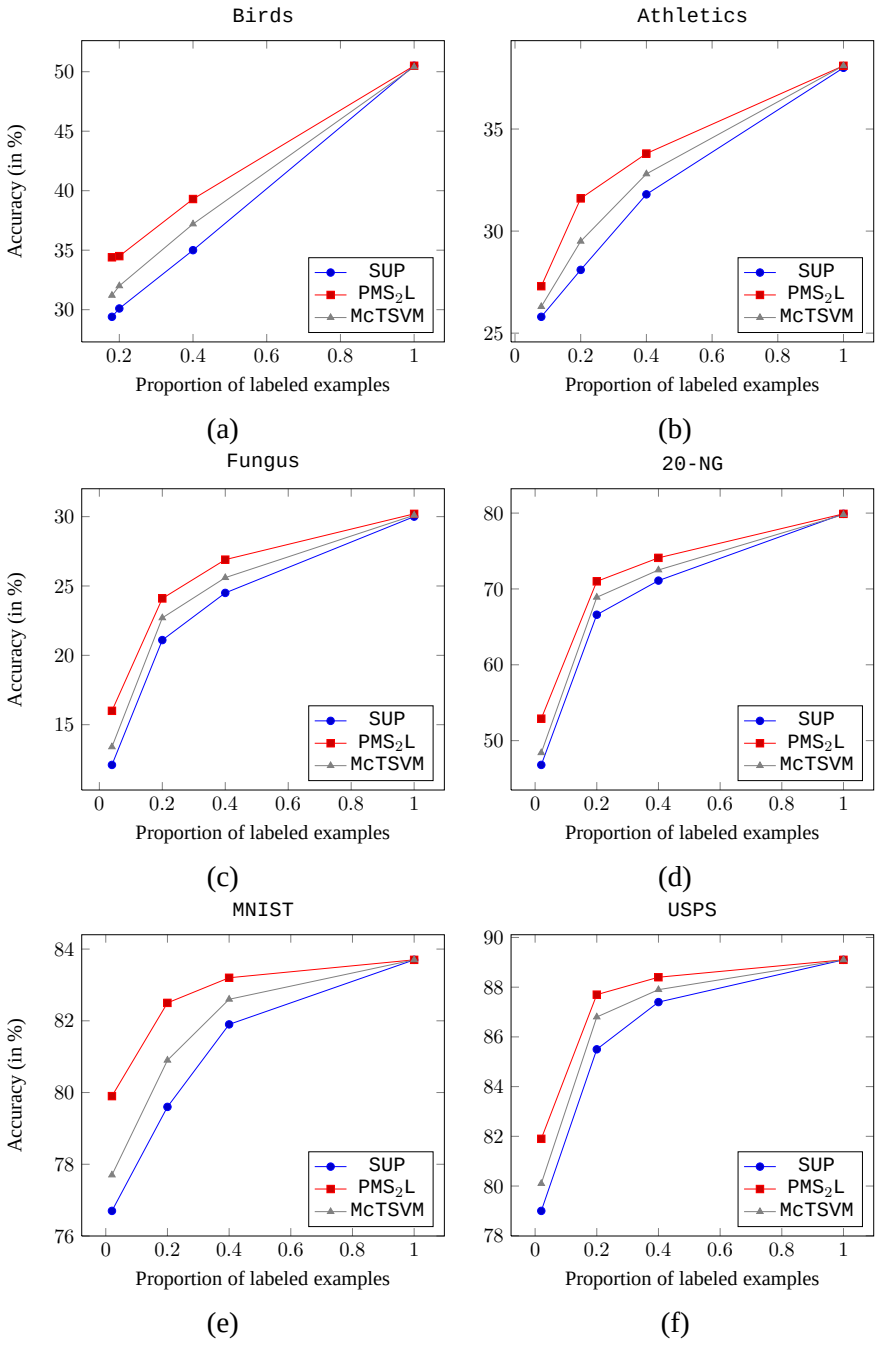

than when only labeled training data are used. Hence, for unlabeled examples within a given cluster, the constraint of predicting the same classes than the $\kappa = 2$ most predominant classes of that cluster makes the decision boundary to pass through regions where the unsigned margins of unlabeled examples are small. As stated in section 4, this is exactly how TSVM works, and the proximity of results between McTSVM and PMS$_2$L, compared to the two other SSL algorithms can be explained by the similitude of the assumptions leading to the development of these models.

However, the fundamental difference between these two algorithms in the iterative pseudo-labeling of unlabeled examples (or not), would do that, when the proportion of labeled training data is small, the iterative pseudo-labeling steps of McTSVM injects noise into the learning process at the same level or even more than the true labeled information. The question therefore arises as to how these two techniques behave for more labeled training data available at the learning phase?

In order to analyze more finely this situation, we compared SUP, PMS$_2$L and McTSVM for an increasing size of the labeled training data. Figure 2, illustrates this by showing the accuracy (in percentage) with respect to the number of labeled examples in the initial labeled training set S_ℓ. The main observations drawn from these results, are:

- As expected, all performance curves increase monotonously with respect to the additional labeled data and converge to the same performance. We note that when all the labeled training data are used for learning the linear SVM gives the same results than those reported in the state-of-the art (e.g. the MLP model with no hidden layer on USPS (LeCun, Bottou, Bengio, & Haffner, 2001) and (Maji & Malik, 2009)).

- Though McTSVM takes advantage of unlabeled data in its learning process, it is outperformed by PMS$_2$L.

- On ImageNet Birds and MNIST, a non-negligible quantity of labeled examples is necessary for SUP to catch the performance of PMS$_2$L learned with the same proportion of labeled data than the one of Table 3, and the remaining unlabeled training data.

These behaviour first suggest that when enough labeled data is available, unlabeled data do not serve the learning algorithm as for the reverse situation. These results suggest that for SSL discriminant techniques designed following the low density separation hypothesis, a more convenient approach than the pseudo-labeling strategy, used in most of these techniques, would be the incorporation of a penalized factor concerning unlabeled examples into the objective of the learning algorithm as the one proposed in Equation 4.

6. Conclusion

The contributions of this paper are twofold. First, we proposed a bound on the risk of a multiclass classifier trained over partially labeled training data. We derived data-dependent bounds for the generalization error of a classifier trained by minimizing an objective function that consists of an empirical risk term, estimated over the labeled training set, and a penalized term corresponding to the ratio of unlabeled examples of each cluster; within the κ bounded set of clusters, for which their predicted class does not belong to the set of the associated κ predominant classes. The analysis of this bound for kernel-based hypotheses reveals a convergence rate that is an extension to the multiclass case, of some other rates over the bounds of the excess risk proposed in the literature. Empirical results on a various datasets support our findings by showing that the proposed algorithm

is competitive compared to different extensions of binary semi-supervised learning algorithms and that it may significantly increase classification performance in the most interesting situation, when there are few labeled data available for training.

References

Amini, M., Truong, T., & Goutte, C. (2008). A boosting algorithm for learning bipartite ranking functions with partially labeled data. In *Proceedings of the 31st Annual International ACM SIGIR Conference*, pp. 99–106.

Balcan, M., & Blum, A. (2010). A discriminative model for semi-supervised learning. *J. ACM*, *57*(3).

Bartlett, P. L., & Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, *3*, 463–482.

Ben-David, S., Lu, T., & Pál, D. (2008). Does unlabeled data provably help? worst-case analysis of the sample complexity of semi-supervised learning. In *21st Annual Conference on Learning Theory - COLT 2008, Helsinki, Finland, July 9-12, 2008*, pp. 33–44.

Blum, A., & Mitchell, T. M. (1998). Combining labeled and unlabeled sata with co-training. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT 1998, Madison, Wisconsin, USA, July 24-26, 1998.*, pp. 92–100.

Bubeck, S., & Luxburg, U. V. (2009). Nearest neighbor clustering: A baseline method for consistent clustering with arbitrary objective functions. *Journal of Machine Learning Research*, *10*, 657–698.

Castelli, V., & Cover, T. M. (1995). On the exponential value of labeled samples. *Pattern Recognition Letters*, *16*(1), 105–111.

Chapelle, O., Schölkopf, B., Zien, A., et al. (2006). *Semi-supervised learning*. MIT press.

El-Yaniv, R., & Pechyony, D. (2009). Transductive rademacher complexity and its applications. *Journal of Artificial Intelligence Research (JAIR)*, *35*, 193–234.

Harchaoui, Z., Douze, M., Paulin, M., Dudík, M., & Malick, J. (2012). Large-scale image classification with trace-norm regularization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, pp. 3386–3393.

Kääriäinen, M. (2005). Generalization error bounds using unlabeled data. In *Learning Theory, 18th Annual Conference on Learning Theory, COLT 2005, Bertinoro, Italy, June 27-30, 2005, Proceedings*, pp. 127–142.

Knuth, D. E. (1976). Big omicron and big omega and big theta. *SIGACT News*, *8*(2), 18–24.

Koltchinskii, V., & Panchenko, D. (2002). Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 1–50.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (2001). Gradient-based learning applied to document recognition. In Haykin, S., & Kosko, B. (Eds.), *Intelligent Signal Processing*, pp. 306–351. IEEE Press.

Ledoux, M., & Talagrand, M. (1991). *Probability in Banach Spaces: Isoperimetry and Processes*, Vol. 23. Springer Science & Business Media.

Lehmann, E. (1975). *Nonparametric Statistical Methods Based on Ranks.* McGraw-Hill, New York.

Lei, Y., Dogan, U., Binder, A., & Kloft, M. (2015). Multi-class svms: From tighter data-dependent generalization bounds to novel algorithms. In *Advances in Neural Information Processing Systems*, pp. 2035–2043.

Leskes, B. (2005). The value of agreement, a new boosting algorithm. In *Learning Theory, 18th Annual Conference on Learning Theory, COLT 2005, Bertinoro, Italy, June 27-30, 2005, Proceedings*, pp. 95–110.

Luxburg, U. V. (2010). Clustering stability: An overview. *Journal Foundations and Trends in Machine Learning*, *2*(3), 235–274.

Luxburg, U. V., Bousquet, O., & Belkin, M. (2004). On the convergence of spectral clustering on random samples: The normalized case. In *Learning Theory, 17th Annual Conference on Learning Theory, COLT 2004, Banff, Canada, July 1-4, 2004, Proceedings*, pp. 457–471.

Maji, S., & Malik, J. (2009). Fast and accurate digit classification. Tech. rep. UCB/EECS-2009-159, EECS Department, University of California, Berkeley.

McDiarmid, C. (1989). On the method of bounded differences. In *Surveys in Combinatorics*, No. 141, pp. 148–188.

Mohri, M., Rostamizadeh, A., & Talwalkar, A. (2012). *Foundations of Machine Learning*. Adaptive computation and machine learning. MIT Press.

Rakhlin, A., & Caponnetto, A. (2006). Stability of k-means clustering. In *Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006*, pp. 1121–1128.

Rigollet, P. (2007). Generalization error bounds in semi-supervised classification under the cluster assumption. *Journal of Machine Learning Research*, *8*, 1369–1392.

Seeger, M. (2000). Learning with labeled and unlabeled data. Tech. rep., Institute for Adaptive and Neural Computation, University of Edinburgh.

Shamir, O., & Tishby, N. (2007). Cluster stability for finite samples. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, pp. 1297–1304.

Shamir, O., & Tishby, N. (2008). On the reliability of clustering stability in the large sample regime. In *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, pp. 1465–1472.

Thiagarajan, J. J., Ramamurthy, K. N., & Spanias, A. (2011). Optimality and stability of the k-hyperline clustering algorithm. *Pattern Recognition Letters*, *32*(9), 1299–1304.

Tolstikhin, I. O., Zhivotovskiy, N., & Blanchard, G. (2015). Permutational rademacher complexity - A new complexity measure for transductive learning. In *Algorithmic Learning Theory - 26th International Conference, ALT*, pp. 209–223.

Urner, R., Shalev-Shwartz, S., & Ben-David, S. (2011). Access to unlabeled data can speed up prediction time. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pp. 641–648.

Vapnik, V. (2000). *The nature of statistical learning theory*. Springer Science & Business Media.

Von Luxburg, U., Belkin, M., & Bousquet, O. (2008). Consistency of spectral clustering. *The Annals of Statistics*, 555–586.

Wang, B., Tu, Z., & Tsotsos, J. K. (2013). Dynamic label propagation for semi-supervised multi-class multi-label classification. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pp. 425–432.

Zhou, D., Bousquet, O., Lal, T. N., Weston, J., & Schölkopf, B. (2003). Learning with local and global consistency. In *Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada]*, pp. 321–328.

Zhu, X. (2002). Learning from labeled and unlabeled data with label propagation. Tech. rep. CMU-CALD-02-107, Carnegie Mellon University.

Zhu, X. (2005). Semi-supervised learning literature survey. technical report 1530. Tech. rep., Department of Computer Sciences, University of Wisconsin.